

- 1
- 2
- 3
- 4
- 5
- 6
- 7
- 8
- 9
- 10
- 11
- 12
- 13
- 14
- 15
- 16
- 17
- 18
- 19
- 20
- 21
- 22
- 23
- 24
- 25
- 26
- 27
- 28

Inventors

Related Applications

Technical Field

Background Art

Databases are perhaps the most critical resource of an enterprise, and therefore it is of prime importance to secure them. Standard database products have access control interfaces for setting permissions on the defined tables and columns of the database. Such access control interfaces can allow, deny, or

1 revoke permissions for a given user and for a given operation on
2 each database table and column.

3 Databases are managed by database administrators (DBAs), who
4 determine access control settings for the database by
5 participating in the design and implementation of the
6 applications using the database and/or by reviewing suggestions
7 made by the application developer or provider. The DBA obtains
8 knowledge of users, groups, roles, and applications accessing the
9 database. This information is not always readily and easily
10 available.
11

12 In any case, the access control settings are configured on
13 the basis of perceived application behavior, i.e., the
14 application is expected to access specific parts of the database
15 for specific operations, and as such, the DBA chooses to apply
16 relevant access control settings. This process is complicated,
17 since database applications are typically huge, with many
18 components developed by large teams. Thus, to consolidate all
19 the features to focus access to parts of the database for
20 specific operations is almost impossible. As a result, it is
21 often the case that unnecessarily loose (open) access control
22 settings are applied to the database to account for the various
23 unknowns.
24
25

26 These loose access control settings give rise to concerns
27 that malicious individuals are thereby able to access parts of
28

1 the database that could have been protected without disturbing
2 access by benign individuals (or applications). They are not
3 prevented from doing so by conventional techniques, since said
4 access is within the bounds of permissible access settings as
5 originally configured by the DBA.
6

7 What is needed is a solution for ensuring that minimal
8 access control settings are applied to the database, so that each
9 application can continue to function as usual, while avoiding
10 loose settings.

11 Disclosure of Invention

12 Computer implemented methods, apparatus, and computer-
13 readable media for empirically adjusting access to a database
14 (1). An apparatus embodiment comprises: coupled to the database
15 (1), a database discovery module (11) for determining authorized
16 accesses to the database (1); coupled to the database (1), a
17 command monitoring module (12) for monitoring actual accesses to
18 the database (1); and coupled to the database discovery module
19 (11) and to the command monitoring module (12), an analysis
20 module (13) for comparing actual accesses with authorized
21 accesses.
22
23

24 Brief Description of the Drawings

25 These and other more detailed and specific objects and
26 features of the present invention are more fully disclosed in the
27
28

1 following specification, reference being had to the accompanying
2 drawings, in which:

3 Figure 1 is a block diagram illustrating apparatus modules
4 suitable for use in the present invention.

5 Figure 2 is a flow diagram illustrating a method embodiment
6 of the present invention.

7 Detailed Description of the Preferred Embodiments

8 "Database" is used broadly herein to comprise any collection
9 of data stored on any computer readable medium. A database
10 normally comprises tables and columns, and is accessed by some
11 query language such as SQL (Structured Query Language).
12

13 "Coupled" is used broadly herein to encompass any type of
14 direct or indirect communicative coupling.
15

16 Figure 1 illustrates apparatus suitable for carrying out the
17 present invention. Database 1 can be any type of database, such
18 as a relational database or a flat file. When database 1 is a
19 relational database, commands 15 are typically written in a SQL
20 language. As used herein, "SQL" is taken in the broad sense to
21 mean the original language known as SQL (Structured Query
22 Language, which originated in the IBM Research Labs in 1969), any
23 derivative thereof, or any structured query language used for
24 accessing a relational database.
25

26 SQL 92 is the current standard version of SQL. It is
27 published in many places, including on the World Wide Web. Each
28

1 vendor of a database 1 tends to have its own flavor of SQL, a
2 flavor compatible with SQL 92.

3 An example of a SQL command 15 is the SELECT command:

4 SELECT C1,C2 FROM T1,T2

5 In the above command, a user is attempting to select columns
6 C1 and C2 from T1 and T2. Each of T1 and T2 can be a table or a
7 view. (Views are discussed below.) Without additional
8 information, we don't know whether C1 belongs to T1 or T2; and we
9 don't know whether C2 belongs to T1 or T2.
10

11 GRANT, DENY, and REVOKE commands can be used to limit access
12 to database 1. Not every database vendor has a REVOKE command.
13 For example, Oracle does not. In this case, DENY is used
14 instead. Most access control modules 16 work on the basis that a
15 DENY command overrides a GRANT command.
16

17 In the case where database 1 is not a relational database,
18 the commands can be written in a language other than SQL, such as
19 XML.

20 Database 1 may have associated therewith an internal audit
21 table 10 and/or an external database log file 19 for storing
22 audit and/or ancillary information pertaining to database 1.
23 Database 1 is typically packaged within a dedicated computer
24 known as a database server, which may also contain communications
25 and other modules. The database server can contain more than one
26 database 1.
27
28

1 Access control module 16 is a hardware, firmware, or
2 software device that determines what users and operations can
3 access the various tables and columns within database 1. Access
4 control settings within module 16 can be set and/or changed by a
5 human database administrator (DBA) 17.

6
7 Database discovery module 11 and command monitoring module
8 12 are coupled to database 1. The purpose of database discovery
9 module 11 is to i.) determine database 1 structure in terms of
10 its tables and columns, and in terms of artifices (views, stored
11 procedures, etc.) that manipulate the tables and columns, and
12 ii.) determine (uncover) authorized (permitted) accesses to
13 database 1. The purpose of command monitoring module 12 is to
14 monitor actual accesses to database 1. Analysis module 13 is
15 coupled to modules 11 and 12. The purpose of analysis module 13
16 is to compare actual database accesses with authorized accesses,
17 and to make and implement appropriate decisions based upon the
18 results of such comparisons. Module 13 can generate one or more
19 third party reports 18. Storage area 14 is coupled to modules 12
20 and 13. The purpose of storage area 14 is to accumulate data
21 generated by command monitoring module 12 during a training phase
22 (observing step 22). Storage area 14 may be part of monitoring
23 module 12.

24
25
26 Modules 11-14 can be implemented in hardware, firmware,
27 software, or any combination thereof. When implemented in
28

1 software, modules 11-14 can reside on one or more computer-
2 readable media, such as one or more hard disks, floppy disks,
3 CDs, DVDs, etc.

4 Command monitoring module 12 is usually a sniffer, because a
5 sniffer does not modify the input stream of data emanating from
6 commands 15. Alternatively, monitoring module 12 can be a proxy.
7 A proxy does affect the input stream, because the queries it
8 receives must be rerouted to the database server. The inputs to
9 module 12 are a plurality of commands 15. Figure 1 illustrates n
10 such commands, where n is a positive integer. Module 12 examines
11 the data streams generated by commands 15 and extracts the
12 relevant information therefrom. Module 12 decrypts commands 15
13 if commands 15 are encrypted.

14 A method embodiment of the present invention will now be
15 described in conjunction with Figure 2.

16 At step 21, authorized accesses of database 1 are discovered
17 by database discovery module 11. An "authorized access"
18 comprises the following combination: the authorized database 1
19 (in those embodiments where the database server has more than one
20 database), the authorized table, the authorized column, the
21 authorized operation, and the authorized user. An individual
22 record produced by database discovery module 11 can have the
23 forms:

24 Permitted[database][table][operation][user]=0 or 1
25
26
27
28

1 Permitted[database][table][column][operation][user]=0 or 1
2 INSERT and DELETE permission settings are made only at the
3 table level, not the column level. One cannot "insert" or
4 "delete" an individual column as this is really an UPDATE
5 operation. So, we will be granting/denying INSERT and DELETE
6 permissions only at the table level. SELECT and UPDATE
7 permission settings can be made at a column granularity, or at a
8 table granularity. Typically, we will be granting/denying SELECT
9 and UPDATE at the column level on a per user basis, unless all of
10 the columns of a given table are granted/denied for a given user
11 for a SELECT or UPDATE, in which case we will do the
12 granting/denying at the table level.

13
14 A view V is a command 15 involving multiple tables. An
15 example of a view V is:

16
17 CREATE V AS SELECT C1,C2 FROM T1,T2

18 The discovery step 21 ascertains the definitions of the
19 views associated with database 1, i.e., discovery module 11
20 resolves the views into individual accessed tables and columns
21 along with the operations used on them. Thus, for example, if a
22 command 15 monitored during the observing step 22 is "SELECT A
23 from V", it is known what columns and tables within database 1
24 are implicated because discovery module 11 has determined this
25 information.

26
27
28

1 A stored procedure is a procedure that is performed on the
2 data within database 1. An example of a stored procedure SP1 is:

3 CREATE SP1 (X,Y) AS STORED PROCEDURE

4 In the above example, X and Y are parameters representing
5 users, commands, tables, columns, dates, functions, constants,
6 etc. Again, database discovery module 11 resolves (breaks down)
7 each stored procedure into accessed tables and columns along with
8 the operations used on them.

9
10 The database discovery module 11 must resolve (break down)
11 any database 1 artifact that performs operations on database
12 tables or columns in a "black box" manner, as do views, stored
13 procedures, user-defined functions, triggers, etc.

14
15 In addition to views and stored procedures, user-defined
16 functions and triggers will also be "discovered". In fact, any
17 SQL entity that reads/writes column values or table rows must be
18 discovered. For example, min/max column value checking can be
19 specified using values from other columns. Also, there can be
20 many levels of nesting. T1 (from the above SELECT example) can
21 really be a view (V). But V itself might be composed of other
22 views (and tables), ad infinitum.

23
24 Module 11 can be programmed to automatically determine the
25 relevant information. Alternatively, or in addition to
26 information determined automatically by module 11, a human such
27 as DBA 17 can provide information to analysis module 13 regarding
28

1 authorized combinations. This can be done, for example, if there
2 are complicated views, stored procedures, etc., associated with
3 database 1.

4 The information uncovered by database discovery module 11 is
5 stored in a storage area, for later use by command monitoring
6 module 12. The storage area can be part of database discovery
7 module 11, or a separate module.

9 During the observing step (training phase) 22, command
10 monitoring module 12 monitors incoming commands 15, keeping track
11 (by means of updating storage area 14) as to which users perform
12 which operations on which tables and columns within database 1.
13 There are four main operations that are tracked. SELECT is used
14 to read a column value. UPDATE is used to modify (write) a
15 column value. INSERT is used to insert a new row of column
16 values into a table. DELETE is used to delete an existing row of
17 column values from a table.

19 Monitoring module 12 first breaks down commands 15 into
20 their constituent elements using information provided by database
21 discovery module 11. For example, in the illustration given
22 above, the command SELECT C1,C2 from T1,T2 is ambiguous in that
23 it is not known a priori whether C1 belongs to T1 or T2, it is
24 not known whether C2 belongs to T1 or T2, and it is not known
25 whether T1 and T2 are tables or views. But this information is
26 determined by database discovery module 11 in step 21, and
27
28

1 command monitoring module 12 uses this information in step 22 to
2 resolve the command and then store the observed combination in
3 storage area 14. An entry in storage area 14 may have one of the
4 forms:

5 Observed[database][table][operation][user]= 0 or 1

6 Observed[database][table][column][operation][user]=0 or 1

7
8 The duration of the observing step 22 is normally defined in
9 terms of a preselected time period. Alternatively, the duration
10 of observing step 21 can be defined in terms of a preselected
11 number of entries made to storage area 14. In either case, the
12 duration of the observing step 22 should be sufficiently long
13 that monitoring module resolves and records many commands 15, and
14 all expected functionalities of the applications accessing
15 database 1 are exercised. This will enable analysis module 13 to
16 completely understand database access patterns.
17

18 Monitoring module 12 can employ any technique of in-line
19 interception or real-time auditing to obtain the desired
20 information.

21 Real-time auditing can be used in cases where database 1 has
22 an auditing feature. The auditing information may be placed into
23 an audit table 10 internal to database 1 or into an external
24 database log file 19. In real-time auditing, module 12 instructs
25 database 1 to generate a stream of events every time a command 15
26 enters database 1. The stream can include such items as the text
27
28

1 of the command 15, a date/time stamp, information pertaining to
2 the user that issued the command 15, the IP (Internet Protocol)
3 address of the issuing computer, the application that issued the
4 command 15, etc. The stream can appear to module 12 in string or
5 binary form, and can be extracted using a number of different
6 techniques, depending upon the implementation, including APIs
7 (Application Programming Interfaces) that access database 1. One
8 example is to use ODBC (Open DataBase Connectivity), a set of C
9 language API's that allows one to examine or modify data within
10 database 1. If the Java programming language is used, JDBC (Java
11 DataBase Connectivity) can be used instead.

12
13 Another way that module 12 extracts the needed information
14 from database 1 is to use code injection or patching to inject
15 logic into one or more modules associated with database 1, to
16 transfer control to module 12.

17
18 In another embodiment, called "direct database integration",
19 the database 1 vendor, who has access to the commands 15 in
20 conjunction with the normal operation of the database 1, makes
21 the commands 15 available to module 12.

22
23 In yet another embodiment, in cases where database 1
24 supports it, external database log file 19 may be examined
25 without the need to resort to special software.

1 Once a command 15 has been processed by module 12, the
2 command 15 can optionally be expunged from any table 10 or log
3 file 19 it is stored in, to make room for subsequent commands 15.

4 Module 12 normally obtains its information in real time,
5 but, alternatively, the information could be extracted in a non-
6 real time manner, e.g., in those embodiments where audit table 10
7 or log file 19 is used.

9 At step 23, analysis module 13 compares the actual accesses
10 of database 1, as gathered in storage area 14, with the normally
11 larger universe of authorized accesses as determined by database
12 discovery module 11. Analysis module 13 can generate a map of
13 which parts (tables and columns) of which database 1 were
14 accessed during step 22 via which operations emanating from which
15 users. The map can then be displayed to DBA 17 by any
16 conventional means. For example, the map can be displayed on a
17 computer monitor, with actual accesses being portrayed in one
18 color, and authorized accesses that were not observed during the
19 observing step 22 portrayed in a different color.

21 At step 24, analysis module 13 and/or DBA 17 adjust access
22 control settings to database 1, based upon results of comparing
23 step 23 and possibly based upon pre-established criteria. The
24 adjustments are made by changing settings within access control
25 module 16. Such adjustments can include one or more of the
26 following:
27
28

1 1) DBA 17 receives a report from module 13 containing
2 suggested revised access control settings that should be applied
3 to database 1. The suggestion is normally to harden database 1,
4 i.e., impose stricter access control settings. For example, the
5 suggestion may be to deny access to operations by certain users
6 on database 1 columns and tables that were preconfigured to be
7 authorized, but which were not observed during observing step 22.

9 2) Analysis module 13 is allowed to automatically harden
10 database 1 (for all times of the day), i.e., apply access control
11 settings as determined by module 13. Again, the criterion can be
12 to deny access to operations by certain users on certain columns
13 and tables of certain databases that were preconfigured to be
14 authorized, but which were not observed during observing step 22.

16 3) Module 13 is allowed to harden database 1 dynamically
17 based upon time of day, i.e., access control module 16 is
18 programmed to harden database 1 during certain times of the day
19 but not during other times.

20 4) DBA 17 can be alerted by module 13 regardless of time of
21 day or in a time-based access pattern. The alerts can convey
22 those combinations of databases, tables, columns, operations, and
23 users that were preconfigured to be authorized, but which were
24 not observed during observing step 22.

26 5) Command monitoring module 12 can be allowed to continue
27 monitoring commands 15 past the duration of the observing step
28

1 22. During the extended time period, for example, analysis
2 module 13 can be programmed to alert DBA 17 in real time
3 regarding commands 15 that are observed by monitoring module 12
4 during the extended time period but that were not observed during
5 observing step 22.

6
7 Adjustments involving the automatic hardening of database 1
8 can be implemented by using a bypass connection 20 between
9 analysis module 13 and access control module 16 to bypass DBA 17.
10 The hardening commands can be written as standard SQL commands as
11 supported by database 1, or by using database specific
12 proprietary APIs (Application Programming Interfaces). Examples
13 of such APIs are OCI for an Oracle database 1, or DMO for a
14 Microsoft SQL Server. OCI and DMO are libraries that enable
15 programming languages other than SQL to access database 1, by
16 translating the commands into SQL. OCI is a set of subroutines
17 in programming language C. DMO uses COM objects in C or in
18 Visual Basic.

19
20 In optional step 25, analysis module 13 generates one or
21 more third party reports 18, e.g., a regulatory compliance report
22 such as the data security report required by HIPAA (Health
23 Insurance Portability and Accountability Act). As used herein,
24 "third party" means a report to be processed by an entity other
25 than DBA 17 or access control module 16.
26
27
28

1 Normal operations of database 1 can continue uninterrupted
2 during steps 21-25 of the present invention. This is an
3 important attribute.

4 Returning to Figure 1, we see an exemplary database 1 having
5 one table (People) with four columns (ID, Name, Phone, and Social
6 Security Number). Command monitoring module 12 monitors two
7 incoming commands 15 from user John. Command 15(1) is an UPDATE,
8 and command 15(2) is a SELECT. Module 12 monitors these commands
9 15 and dissects which operations are being performed on which
10 tables and columns within database 1. At step 23, analysis
11 module 13 concludes that user John has not selected column SSN,
12 has not updated column ID, and has not updated column SSN. (John
13 has implicitly selected column ID, because of the WHERE
14 subcommand embedded within command 15(1)). At step 24, module 13
15 can perform one or more tasks as described above, such as to
16 automatically deny future access to John to these unobserved
17 combinations, generate an alert to DBA 17 if a subsequent command
18 15 arrives that attempts to access one of these previously
19 unobserved combinations, etc.

20 The above description is included to illustrate the
21 operation of the preferred embodiments and is not meant to limit
22 the scope of the invention. The scope of the invention is to be
23 limited only by the following claims. From the above discussion,
24 many variations will be apparent to one skilled in the art that
25

1 would yet be encompassed by the spirit and scope of the present
2 invention.

3 What is claimed is:
4

5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28